Differentially Private LQ Control

Kasra Yazdani*, Austin Jones†, Kevin Leahy†, Matthew Hale*

***Abstract*—As multi-agent systems proliferate and share more user data, new approaches are needed to protect sensitive data while still enabling system operation. To address this need, this paper presents a private multi-agent LQ control framework. Agents' state trajectories can be sensitive and we therefore protect them using differential privacy. We quantify the impact of privacy along three dimensions: the amount of information shared under privacy, the control-theoretic cost of privacy, and the tradeoffs between privacy and performance. These analyses are done in conventional control-theoretic terms, which we use to develop guidelines for calibrating privacy as a function of system parameters. Numerical results indicate that system performance remains within desirable ranges, even under strict privacy requirements.**

I. Introduction

MULTI-AGENT systems, such as smart power grids and robotic swarms, require agents to exchange information to work together. In some cases, the information shared may be sensitive. For example, consumption data in a power grid can expose habits and activities of individuals [1], [2]. Sensitive user data must be protected when it is shared, though of course it must remain useful in multi-agent coordination. Hence, privacy in multi-agent control should protect sensitive data from the agent receiving it while still ensuring that private data remains useful to that recipient.

Recently, privacy of this form has been achieved using differential privacy. Differential privacy was originally designed to protect data of individuals in static databases [3], [4]. Its goal is to allow accurate statistical analyses of a population while providing strong, provable privacy guarantees to individuals. Differential privacy is appealing because it is immune to post-processing [5], in that post-hoc computations on private data do not weaken privacy's guarantees. For example, filtering private trajectories can be done without harming privacy [6], [7]. Differential privacy is also robust to side information [8], in that its privacy guarantees are not defeated by an adversary with access to additional information about data-producing entities. Differential privacy has been extended to dynamical systems [6], [9], [10] in which trajectory-valued data is protected, and it is this notion of differential privacy that we use.

Linear-quadratic (LQ) control is the underlying framework for many existing multi-agent control applications. One example is smart power systems where power forecast, generation, and distribution require access to time series usage data measured by smart meters. In particular, LQ control for load frequency control has been used in power systems [11], [12], with the objective of restoring balance between power consumption and generation. In addition, the work in [13]–[16] incorporates an LQ control scheme for stability and performance of wide-area power control systems using standard phasor measurement units. Other applications of LQ control include motion planning [17] and security of cyber physical systems [18]. Existing work investigates convergence and performance under various constraints; however, despite the sensitive nature of the data involved, privacy is generally absent in their treatment.

In this paper, we use differential privacy to develop a private multi-agent LQ control framework. Adding privacy noise makes this problem equivalent to a linear quadratic Gaussian (LQG) problem, and the optimal controller will be linear in the expected value of agents' states. Computing this expected value is a centralized operation, and we therefore augment the network with a cloud computer [19]. In contrast to some existing approaches, the cloud is *not* a trusted third party and does not receive sensitive information from any agent [20]. The cloud instead gathers private information from the agents, estimates their states, and generates optimal inputs. These inputs are transmitted back to the agents, which apply them in their local state updates, and then this process repeats.

Contributions: Although there exists a large body of privacy research, privacy parameter interpretation and selection both largely remain the domain of subject matter experts. Moreover, since offering privacy guarantees for a control system generally involves sacrificing some level of performance, it is critical to quantify the effects of privacy to rigorously evaluate tradeoffs. Our contributions are therefore the following:

1) Developing an algorithm for multi-agent differentially private LQ control. (*Sec. IV*)
2) Quantifying sensitive information revealed by bounding filter accuracy in terms of privacy parameters (*Sec. V*)
3) Providing quantitative criteria for privacy calibration to trade off information shared and control cost (*Sec. VI*)
4) Quantifying the relationship between agents' cost and their privacy levels (*Sec. VII*)

Preliminary versions of this work appeared in [21], [22]. This paper differs from [21] because it does not rely on a trusted aggregator. Further, we quantify the tradeoff between cost and privacy, which was not explored in [21], [22].

*Kasra Yazdani and Matthew Hale are with the Department of Mechanical and Aerospace Engineering at the University of Florida, Gainesville, FL USA. Emails: {kasra.yazdani,matthewhale}@ufl.edu.

†Austin Jones and Kevin Leahy are with the Massachusetts Institute of Technology Lincoln Laboratory, Lexington, MA USA. Emails: {austin.jones,kevin.leahy}@ll.mit.edu.

KY and MH were supported in part by NSF CAREER Grant No. 1943275 and by AFOSR under Grant FA9550-19-1-0169.

DISTRIBUTION STATEMENT A:Approved for public release. Distribution is unlimited.

This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

Organization: Section II reviews privacy background. Section III defines the private LQG problem, and Section IV solves it. Section V bounds filter error under privacy, and in Section VI we provide guidelines for calibrating privacy. Section VII quantifies the cost of privacy. Next, we provide simulations in Section VIII, and then Section IX concludes.

II. Review of Differential Privacy

Differential privacy is a statistical notion of privacy that masks sensitive data while still enabling accurate analyses of it [5]. It is appealing because post-processing does not weaken its protections. In particular, filtering private data is permitted. Moreover, differential privacy is not weakened even if an adversary knows the privacy mechanism used [4], [5]. We briefly review differential privacy here and refer the reader to [4]–[6] for a thorough introduction.

We use the "input perturbation" approach to differential privacy, which means that agents add noise directly to their outputs before sharing them. Thus, agents do not ever share sensitive data. Privacy guarantees are likewise provided on an individual basis. Formally, each agent's state trajectory will be made approximately indistinguishable from other nearby state trajectories which that agent individually could have produced.

We use the notation $[\ell] = \{1, \dots, \ell\}$ for $\ell \in \mathbb{N}$. We consider trajectories of the form $Z = (Z(1), Z(2), \dots)$, where $Z(k) \in \mathbb{R}^d$ and $\|Z(k)\|_2 < \infty$ for all k. Denote the set[1] of all such sequences by ℓ_2^d.

We consider N agents, and we denote agent i's state trajectory by $x_i \in \tilde{\ell}_2^{n_i}$ for some $n_i \in \mathbb{N}$. The k^{th} element of x_i is $x_i(k) \in \mathbb{R}^{n_i}$. We define our adjacency relation over $\tilde{\ell}_2^{n_i}$.

Definition 1. (*Adjacency for Trajectories*) Fix an adjacency parameter $b_i > 0$ for agent i. Two trajectories $v_i, w_i \in \tilde{\ell}_2^{n_i}$ are adjacent if $\|v_i - w_i\|_{\ell_2} \leq b_i$. We write $\mathrm{Adj}_{b_i}(v_i, w_i) = 1$ if v_i, w_i are adjacent, and $\mathrm{Adj}_{b_i}(v_i, w_i) = 0$ otherwise.

This adjacency relation requires that every agent's state trajectory be made approximately indistinguishable from all other state trajectories not more than distance b_i away. Next, we define differential privacy for dynamic systems. This definition considers outputs of agent i of dimension q_i at each point in time. Output signals are in the set $\tilde{\ell}_2^{q_i}$, over which we use the σ-algebra $\Sigma_2^{q_i}$ (see [6] for a formal construction).

Definition 2. *(Differential Privacy for Trajectories)* Let $\epsilon_i > 0$ and $\delta_i \in (0, 1/2)$ be given. A mechanism M is (ϵ_i, δ_i)-differentially private if, for all adjacent $x_i, x_i' \in \tilde{\ell}_2^{n_i}$, we have

$$\mathbb{P}[M(x_i) \in S] \leq e^{\epsilon_i} \mathbb{P}[M(x_i') \in S] + \delta_i \text{ for all } S \in \Sigma_2^{q_i}.$$

We enforce this definition with the Gaussian mechanism, defined next. We use $s_1(\cdot)$ for the largest singular value of a matrix, and $\mathcal{Q}$ to denote the Gaussian tail integral [23].

Lemma 1 (*Gaussian mechanism*; [6])**.** Let agent i use privacy parameters $\epsilon_i > 0$ and $\delta_i \in (0, 1/2)$ and adjacency parameter $b_i > 0$. For outputs $y_i(k) = C_i x_i(k)$, the Gaussian mechanism sets $\tilde{y}_i(k) = y_i(k) + v_i(k)$, with $v_i(k) \sim \mathcal{N}(0, \sigma_i^2 I_{q_i})$, where I_{q_i} is the $q_i \times q_i$ identity matrix, and $\sigma_i \geq \frac{s_1(C_i) b_i}{2\epsilon_i}(K_{\delta_i} + \sqrt{K_{\delta_i}^2 + 2\epsilon_i})$, with $K_{\delta_i} := \mathcal{Q}^{-1}(\delta_i)$. This is (ϵ_i, δ_i)-differentially private with respect to Adj_{b_i}.

For convenience, we set $\kappa(\delta_i, \epsilon_i) = \frac{1}{2\epsilon_i}(K_{\delta_i} + \sqrt{K_{\delta_i}^2 + 2\epsilon_i})$. We use the Gaussian mechanism for the rest of the paper.

[1]This notation comes from the fact that all such Z have finite truncations with finite ℓ_2-norm. See [6] for additional discussion.

III. Problem Formulation

We next introduce the private multi-agent LQG problem. Below, we write $\mathrm{diag}(P_1, \dots, P_n) := \bigoplus_{i=1}^n P_i$ for matrices P_1 through P_n.

A. Multi-Agent LQ Formulation

Consider N agents indexed over $i \in [N]$. At time k, agent i has state $x_i(k) \in \mathbb{R}^{n_i}$, with dynamics

$$x_i(k+1) = A_i x_i(k) + B_i u_i(k) + w_i(k),$$

where $u_i(k) \in \mathbb{R}^{m_i}$, $w_i(k) \in \mathbb{R}^{n_i}$, $A_i \in \mathbb{R}^{n_i \times n_i}$, and $B_i \in \mathbb{R}^{n_i \times m_i}$. The distribution of process noise is $w_i(k) \sim \mathcal{N}(0, W_i)$, where $W_i \in \mathbb{R}^{n_i \times n_i}$ is symmetric and positive definite. All process noise terms are independent.

We define the state $x(k) = (x_1^T(k) \dots x_N^T(k))^T \in \mathbb{R}^n$ and control $u(k) = (u_1^T(k) \dots u_N^T(k))^T \in \mathbb{R}^m$, where the dimensions $n = \sum_{i \in [N]} n_i$ and $m = \sum_{i \in [N]} m_i$. Along with $w(k) = (w_1^T(k), \dots, w_N^T(k))^T \in \mathbb{R}^n$, and the matrices $A = \mathrm{diag}(A_1, \dots, A_N) \in \mathbb{R}^{n \times n}$ and $B = \mathrm{diag}(B_1, \dots, B_N) \in \mathbb{R}^{n \times m}$, we have the dynamics

$$x(k+1) = Ax(k) + Bu(k) + w(k).$$

We consider infinite-horizon problems with cost

$$J(x,u) = \lim_{K_f \to \infty} \frac{1}{K_f} \mathbb{E}\left\{ \sum_{k=1}^{K_f} (x(k) - \bar{x}(k))^T Q (x(k) - \bar{x}(k)) + u(k)^T R u(k) \right\},$$

where $Q \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$. The vector $\bar{x}_i(k) \in \mathbb{R}^{n_i}$ is agent i's desired state at time k, and we define $\bar{x}(k) = (\bar{x}_1^T(k), \dots, \bar{x}_N^T(k))^T$. We make the standard assumption that $\lim_{k \to \infty} x(k) = \bar{x}$ exists [24].

Assumption 1. In the cost J, $Q = Q^T \succ 0$ and $R = R^T \succ 0$. The pair (A, B) is controllable, and there exists Ω such that $Q = \Omega^T \Omega$ and such that the pair (A, Ω) is observable.

Assumption 1 is standard in LQ control [24]–[26], and it guarantees the existence of a solution to an algebraic Riccati equation that we will encounter below [24, Chapter 4].

B. Differentially Private Information Sharing

The cost J is generally non-separable, which means that it cannot be minimized by agents using only knowledge of their own states. We therefore introduce a cloud computer to aggregate information and distribute control inputs to the agents. The cloud has been used in cyber-physical systems, e.g., in SCADA-based monitoring and state estimation [13]–[15], [27], and is a natural choice here.

At time k, the cloud requests from agent i the output value $y_i(k) = C_i x_i(k)$, where $C_i \in \mathbb{R}^{q_i \times n_i}$. To protect its state trajectory, agent i sends a differentially private form of y_i to the cloud. The cloud uses these private outputs to compute optimal inputs for the agents. Agents use these inputs in their local state updates, and then this process repeats.

Agent i adds noise to each $y_i(k)$ before sending it to the cloud to enforce differential privacy for x_i. Agent i selects privacy parameters $\epsilon_i > 0$ and $\delta_i \in (0, 1/2)$ and adjacency parameter $b_i > 0$. Then agent i sends the cloud

$$\tilde{y}_i(k) := y_i(k) + v_i(k) = C_i x_i(k) + v_i(k), \qquad (1)$$

where the privacy[2] noise $v_i(k) \sim \mathcal{N}\left(0, \sigma_i^2 I_{q_i}\right)$ and $\sigma_i \geq \kappa(\delta_i, \epsilon_i)\, s_1(C_i)\, b_i$ from Lemma 1. The full privacy vector is $v(k) = (v_1^T(k), \ldots, v_N^T(k))^T$, and, for all k, we have $v(k) \sim \mathcal{N}(0, V)$ with $V = \text{diag}\left(\sigma_1^2 I_{q_1}, \ldots, \sigma_N^2 I_{q_N}\right)$. Below we use $C = \text{diag}(C_1, \ldots, C_N)$.

Agents' reference trajectories are a source of side information that can reveal their intentions. However, agents do not need to reveal their whole reference trajectories to the cloud. As written, the cost J depends on $\bar{x}(k)$ for all k, but, leveraging the standard average cost-per-stage formulation, one can replace $\bar{x}(k)$ with $\bar{x}$ for all k with no loss of optimality; see [24] for a thorough discussion. We emphasize that this change is independent of privacy and is a standard approach in infinite-horizon LQG. As a result, only the limit of agent i's reference trajectory, denoted $\bar{x}_i$, is needed by the cloud to compute optimal inputs. Agent i thus privatizes $\bar{x}_i$ before sharing it[3]. Agent i selects privacy parameters $\bar{\epsilon}_i > 0$ and $\bar{\delta}_i \in (0, 1/2)$ and adjacency parameter β_i. Two reference limits $\bar{x}_i, \bar{x}_i'$ are adjacent if $\|\bar{x}_i - \bar{x}_i'\|_2 \leq \beta_i$. Then privacy noise is added via $\tilde{x}_i := \bar{x}_i + \bar{w}_i$. Using the rules for privatizing static data in [6, Lemma 1], agent i generates noise via $\bar{w}_i \sim \mathcal{N}\left(0, \bar{\sigma}_i^2 I_{n_i}\right)$, with $\bar{\sigma}_i \geq \kappa\left(\bar{\delta}_i, \bar{\epsilon}_i\right) \beta_i$.

Problem 1. Let the initial estimate $\hat{x}(0) = \mathbb{E}[x(0)]$ and the matrices A, B, C, V, and W be public information. Minimize

$$\tilde{J}(x, u) = \lim_{K_f \to \infty} \frac{1}{K_f} E\left\{ \sum_{k=1}^{K_f} (x(k) - \tilde{x})^T Q (x(k) - \tilde{x}) + u(k)^T R u(k) \right\}$$

over all control signals u with $u(k) \in \mathbb{R}^m$, subject to

$$x(k+1) = Ax(k) + Bu(k) + w(k)$$
$$\tilde{y}(k) = Cx(k) + v(k),$$

where agent i has privacy parameters (ϵ_i, δ_i) and $(\bar{\epsilon}_i, \bar{\delta}_i)$.

[2]In Equation (III-B), measurement noise inherent to the system can be included, and all analyses permit this change. The form of various Riccati equations will remain the same, with instances of V replaced by the sum of V and the measurement noise covariance matrix. However, we focus on bounding the effects of privacy, and thus exclude measurement noise.

[3]We expect privatizing the reference limit to be unproblematic in applications in which it only changes the cost incurred, e.g., in applications where states are non-physical quantities. However, if the reference also encodes some notion of safety that could be affected by privacy, e.g., collision avoidance, the approach we present can be augmented with a low-level reactive controller for that purpose, such as a control barrier function [28].

IV. Private LQG Tracking Control

Problem 1 is an infinite-horizon LQG problem whose optimal controller [29] is

$$u^*(k) = L\hat{x}(k) + Mg, \qquad (2)$$

where $M = -\left(R + B^T K B\right)^{-1} B^T$ and $L = MKA$. Here, K is the unique positive semidefinite solution to the discrete algebraic Riccati equation

$$K = A^T K A - A^T K B \left(R + B^T K B\right)^{-1} B^T K A + Q$$

and g solves $g = A^T[I - KB(R + B^T KB)^{-1}B^T]g - Q\tilde{x}$. Without privacy, g would depend on $\bar{x}$, but the cloud only receives its private form, $\tilde{x}$, and this is what it must use.

Computing state estimates for infinite time horizons can use a time-invariant Kalman filter [24, Section 5.2] whose prediction step is $\hat{x}^-(k+1) = A\hat{x}(k) + Bu(k)$. The *a posteriori* state estimate $\hat{x}(k)$ is computed with

$$\hat{x}(k+1) = \hat{x}^-(k+1) + \overline{\Sigma} C^T V^{-1}\left(\tilde{y}(k+1) - C\hat{x}^-(k+1)\right),$$

where the *a posteriori* error covariance matrix $\overline{\Sigma}$ is given by $\overline{\Sigma} = (C^T V^{-1} C + \Sigma^{-1})^{-1}$, and the *a priori* error covariance Σ is the unique positive semidefinite solution to the discrete algebraic Riccati equation $\Sigma = A(\Sigma^{-1} + C^T V^{-1} C)^{-1} A^T + W$. The terms $K, L, M, \overline{\Sigma}, \Sigma$, and g can be all computed beforehand by the cloud to reduce its computational load at runtime.

We solve Problem 1 in Algorithm 1: for all $i \in [N]$, Algorithm 1 provides (ϵ_i, δ_i)-differential privacy for agent i's state trajectory and $(\bar{\epsilon}_i, \bar{\delta}_i)$-differential privacy for $\bar{x}_i$.

Algorithm 1: Differentially Private LQG (Solution to Problem 1)

Data: Public information: A_i, B_i, C_i, ϵ_i, δ_i, $\bar{\epsilon}_i$, $\bar{\delta}_i$, $\hat{x}_i(0)$, W_i, and V_i for all i, and Q, R

For all i, agent i chooses (ϵ_i, δ_i) and $(\bar{\epsilon}_i, \bar{\delta}_i)$. It computes $\tilde{x}_i$ and sends it to the cloud

In the cloud, compute K, L, M, Σ, $\overline{\Sigma}$, and g

for $k = 0, 1, 2, \ldots$ **do**

  for $i = 1, \ldots, N$ **do**

    Agent i sends the cloud the private output $\tilde{y}_i(k) := C_i x_i(k) + v_i(k)$

  In the cloud, compute $u^*(k)$ via (IV), send $u_i^*(k)$ to agent i

  for $i = 1, \ldots, N$ **do**

    Agent i updates its state via $x_i(k+1) = A_i x_i(k) + B_i u_i^*(k) + w_i(k)$

The feedback control signals u_i^*, $i \in [N]$, are computed using estimates of agents' states, and these state estimates are functions of the private output trajectories $\tilde{y}_i$, $i \in [N]$. The signals u_i^* are thus post-processing on private data and do not reveal agents' state trajectories. In addition, knowledge of how u_i^* depends upon the x_i's is equivalent to knowledge of agents' dynamics, which is often assumed to be public information and is unproblematic for privacy. Therefore, this use of a feedback controller does not harm privacy.

V. Quantifying Error Induced by Privacy

Algorithm 1 solves Problem 1, though adding privacy noise makes it more difficult for the cloud to compute optimal control values. Indeed, the purpose of differential privacy is to protect an agent's state from the cloud, other agents, and any eavesdroppers. Thus, the cloud is forced to estimate agents' states to generate control values for them. Accordingly, in this section we quantify the ability of the cloud to estimate the agents' states as a measure of the impact of privacy.

The cloud runs a Kalman filter and computes the input $u^*(k)$, though privacy noise only affects the Kalman filter due to the certainty equivalence principle [24]. We therefore quantify the impact of privacy upon the Kalman filter in Algorithm 1 by investigating the best estimate that can be computed with differentially private outputs.

We proceed by developing trace bounds for the *a priori* error covariance matrix Σ and the *a posteriori* error covariance matrix $\overline{\Sigma}$, which are, respectively, equal to the steady-state mean-square error (MSE) of the prediction and estimation steps in the Kalman filter. Because the Kalman filter minimizes both of these quantities, lower bounds on them are lower bounds on (asymptotic) MSE for *any* filtering strategy.

We use $\lambda_n(\Upsilon) \leq \cdots \leq \lambda_1(\Upsilon)$ to denote the ordered eigenvalues of the matrix Υ. For simplicity, consider C diagonal. Noting that $C^TV^{-1}C = \text{diag}\left(\frac{C_{11}^2}{\sigma_1^2}, \ldots, \frac{C_{nn}^2}{\sigma_n^2}\right)$, we define

$$l = \underset{1\leq i\leq n}{\arg\min} \frac{C_{ii}^2}{\sigma_i^2}, \qquad u = \underset{1\leq i\leq n}{\arg\max} \frac{C_{ii}^2}{\sigma_i^2}. \tag{3}$$

Theorem 1. Suppose every agent shares its private output trajectory, and the cloud has all public information. Then the steady-state *a priori* MSE of the Kalman filter is bounded via

$$\text{tr}W + \frac{\sigma_u^2\text{tr}(A^TA)\lambda_n(W)}{\sigma_u^2 + \lambda_n(W)C_u^2} \leq \text{tr}\Sigma \leq \text{tr}W + \frac{\sigma_l^2\text{tr}(A^TA)}{C_l^2}$$

and the steady-state *a posteriori* MSE is bounded via

$$\frac{n\sigma_u^2}{C_u^2 + \sigma_u^2\lambda_n^{-1}(W)} \leq \text{tr}\overline{\Sigma} \leq n\frac{\sigma_l^2}{C_l^2},$$

where $\sigma_l = \kappa(\delta_l, \epsilon_l)s_1(C_l)b_l$ and $\sigma_u = \kappa(\delta_u, \epsilon_u)s_1(C_u)b_u$ are the minimum and maximum privacy noise among agents.

Proof: See the appendix. ∎

These bounds relate privacy to the accuracy of information shared with the cloud and give insight into differential privacy's protections in conventional estimation-theoretic terms. We next leverage these bounds to guide privacy calibration.

VI. Guidelines for Selecting Privacy Parameters

Calibrating privacy can be challenging. The computer science literature has studied this problem [30], though, to the best of our knowledge, there are not control-theoretic guidelines for calibrating privacy. Therefore, in this section, we develop such techniques. The privacy parameter δ can be interpreted as the probability that ϵ-differential privacy fails, and is typically [6] chosen in the range $[10^{-5}, 10^{-1}]$ on this basis. The parameter ϵ can be interpreted as the privacy loss of differential privacy, and it is typically the parameter to be tuned. We therefore develop guidelines for calibrating ϵ.

Theorem 2. Suppose the cloud has all public information, and agent i shares its private output trajectory $\tilde{y}_i$, where $\tilde{y}_i(k) = C_ix_i(k) + v_i(k)$. Take $\delta_i \in [10^{-5}, 10^{-1}]$ and set $\sigma_i = s_1(C_i)\kappa(\delta_i, \epsilon_i)b_i$. Suppose we want the MSE in the cloud's state estimates to be bounded below by $B_l > 0$ and above by $B_u > B_l$. These bounds are attained if

$$\frac{1}{8}\left(\frac{1+\sqrt{36\eta_4+1}}{\eta_4}\right)^2 \leq \epsilon_i \leq \frac{1}{\eta_3}$$

for all i, where

$$\eta_3 := \left(\frac{B_lC_u^2}{s_1(C_i)^2b_i^2\left(n - B_l\lambda_n^{-1}(W)\right)}\right)^{1/2}, \; \eta_4 := \left(\frac{B_uC_l^2}{ns_1(C_i)^2b_i^2}\right)^{1/2}.$$

Proof: See the appendix. ∎

Theorem 2 provides guidelines for choosing ϵ_i, which allows agents to make informed decisions for privacy. With this ability, we next examine privacy's impact upon the cost J.

VII. The Control-Theoretic Cost of Privacy

Implementing differential privacy adds noise where it would otherwise be absent, and we expect privacy to increase the cost J relative to a non-private implementation. Without any cost considerations, one could add noise of very large variance to provide arbitrarily strong privacy. However, private information is used to compute control inputs, which affect future states. Thus, there is a need to balance privacy and performance. The existing literature has explored several notions of a "cost of privacy;" LQG minimizes J, and we therefore compute the increase in J due to privacy, which offers a "cost of privacy" in standard control-theoretic terms.

Theorem 3 (*Cost of Privacy*)**.** Let $J_0(x,u)$ be the cost of Algorithm 1 without privacy, i.e., with $v_i(k) = \bar{w}_i = 0$ for all i and k. Let $\tilde{J}(x,u)$ be the cost of Algorithm 1 with privacy. Then the cost of privacy in LQG, denoted ΔJ, is

$$\begin{aligned}\Delta J(x,u) &= \tilde{J}(x,u) - J_0(x,u) \\ &= \text{tr}\left(K\Sigma + (Q-K)\,\overline{\Sigma}\right) - \text{tr}\,(KW) \\ &\quad + \text{tr}(Q\overline{W}) + \text{tr}(H^TRH\overline{W}),\end{aligned} \tag{4}$$

where $H = M\left[I - (A+BL)^T\right]^{-1}$.

Proof: See the appendix. ∎

After selecting a privacy level and computing its cost, agents may wish to change their privacy levels to tune costs. For example, agents may choose to relax privacy for significant reductions in ΔJ. A natural way to analyze these changes is with the derivative of the cost of privacy ΔJ with respect to ϵ_i; recalling that δ_i is typically fixed *a priori*, ϵ_i is the parameter to be tuned. For simplicity, we take $\bar{\epsilon}_i = \epsilon_i = \epsilon$, $\bar{\delta}_i = \delta_i = \delta$, and $s_1(C_i)b_i = \omega$ for all i.

Theorem 4. Let A be stable. Then the sensitivity of the cost of privacy to changes in privacy is lower-bounded via

$$\frac{d\Delta J}{d\epsilon} \geq \left(-\frac{\omega}{\epsilon}\kappa(\delta,\epsilon)+\frac{\omega}{2\epsilon}\frac{1}{\sqrt{K_\delta^2+2\epsilon}}\right)\cdot \left(\lambda_1(K)\frac{-2\sigma\mathrm{tr}(F^TF)}{\lambda_n(U)}+2\sigma\mathrm{tr}Q+2\sigma\mathrm{tr}\left(H^TRH\right)+\lambda_1(Q-K)\left[\max\left\{\frac{-2\sigma\mathrm{tr}(F^TF)}{\lambda_1(U)},0\right\}\lambda_n\left(\bar{U}\right)+\mathrm{tr}\left(2\sigma\bar{F}^T\bar{F}\right)\right]\right)$$

and upper-bounded via

$$\frac{d\Delta J}{d\epsilon} \leq \left(-\frac{\omega}{\epsilon}\kappa(\delta,\epsilon)+\frac{\omega}{2\epsilon}\frac{1}{\sqrt{K_\delta^2+2\epsilon}}\right)\cdot \left(\lambda_n(K)\max\left\{\frac{-2\sigma\mathrm{tr}(F^TF)}{\lambda_1(U)},0\right\}+2\sigma\mathrm{tr}Q+2\sigma\mathrm{tr}\left(H^TRH\right)\lambda_n(Q-K)\left[\frac{-2\sigma\mathrm{tr}(F^TF)}{\lambda_n(U)}\lambda_1\left(\bar{U}\right)+\mathrm{tr}\left(2\sigma\bar{F}^T\bar{F}\right)\right]\right),$$

where we use the matrices $P=C^T\left(C\Sigma C^T+V\right)^{-1}C\Sigma A^T$, $U=(A^T-P)(A-P^T)-I$, $\bar{P}=C^T\left(C\Sigma C^T+V\right)^{-1}C\Sigma$, $F=\left(C\Sigma C^T+V\right)^{-1}C\Sigma A^T$, $\bar{F}=\left(C\Sigma C^T+V\right)^{-1}C\Sigma$, and $\bar{U}=(I-\bar{P})(I-\bar{P}^T)$.

Proof: See the appendix. ■

Theorem 4 explores the continuum of privacy costs that result from varying ϵ, and for a given problem it provides parameter regimes that either make it useful or not to relax privacy. One may also wish to enforce hard constraints on performance. Next, we provide guidelines for choosing the privacy parameters $\{\epsilon_i\}_{i\in[N]}$ to enforce a desired cost bound.

Theorem 5. Suppose a performance requirement is given as a bound on cost by requiring $\tilde{J}(x,u) \leq \alpha$. Take $\delta_i \in [10^{-5},10^{-1}]$ and set $\sigma_i = s_1(C_i)\kappa(\delta_i,\epsilon_i)b_i$. Then Algorithm 1 attains $\tilde{J}(x,u) \leq \alpha$ if, for all i, $\epsilon_i \geq \frac{1}{8}\left(\frac{1+\sqrt{36\eta_5+1}}{\eta_5}\right)^2$, where

$$\eta_5 = \left[\frac{B_u-\lambda_1(K)\mathrm{tr}W-\bar{x}^TQ\bar{x}+g^TB\left(R+B^TKB\right)^{-1}B^Tg}{s_1(C_i)^2b_i^2\left(\lambda_1(K)\frac{\mathrm{tr}(A^TA)}{C_l^2}+\mathrm{tr}(H^TRH)+\mathrm{tr}(Q)\right)}\right]^{1/2}.$$

Proof: See the appendix. ■

VIII. Case Study

Load Frequency Control (LFC) regulates power flow to different areas while balancing load and generation. In our framework, each area is an agent, and we consider a system of ten decoupled areas. LFC requires transmitting measurements from remote terminal units (RTUs) to a control center and control signals from the control center to the plant side. This aggregation and communication have well-established privacy concerns [1], [2], and we use Algorithm 1 for it. The continuous time dynamic model of the multi-area LFC system

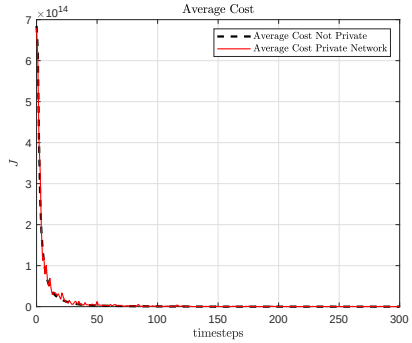

Fig. 1: The time-average cost incurred by ten power system areas with privacy (solid line) and without privacy (dashed line). Privacy increases costs, though these increases become small relative to the noise-free cost, indicating that privacy's impact is not excessive. This simulation was performed 100 times and the average over these runs is plotted.

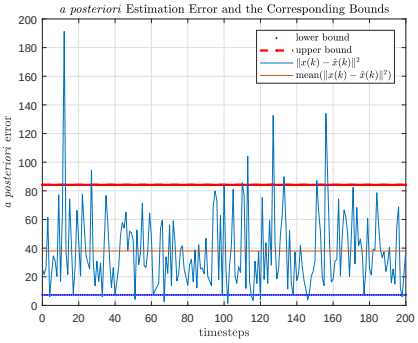

Fig. 2: The squared error of the cloud's state estimates (solid line), the lower bound on estimation error in Theorem 1 (dotted line), and the upper bound on estimation error in Theorem 1 (dashed line) over 200 timesteps. We see that instantaneous *a posteriori* error typically obeys our MSE bounds, and on average lies within the bounds.

is given by $\dot{X}(t)=A_cX(t)+B_cU(t)$, and the matrices A_c and B_c can be found in [12]. The state vector for agent i is

$$x_i(t)=[\Delta f^i(t),\ \Delta P_g^i(t),\ \Delta P_{tu}^i(t),\ \Lambda^i(t)]^T\in\mathbb{R}^4,$$

where $\Delta f^i(t), \Delta P_g^i(t)$, and $\Delta P_{tu}^i(t)$ are the frequency deviation, generator power deviation, and position value of the turbine, respectively. The control input error on the i-th power area is denoted by $\Lambda^i(t)=\int_0^t\vartheta_i\Delta f^i(s)dt$, where ϑ_i is the frequency bias factor. We simulate 5 agents with dynamics of Area 1 from [12] and 5 agents with dynamics of their Area 2.

We discretize the dynamics of $X(t)$ with $A=e^{A_ch}$ and $B=\int_0^h e^{A_c\tau}B_cd\tau$, where h is the sampling period. We have $C=I_{40\times40}$ and $W=I_{40\times40}$. We initialize all states to zero. All areas select identical privacy parameters, namely, $(\epsilon_i,\delta_i)=(\ln 3, 0.001)$ for all i. In addition, $\bar{x}_i=0$ is made private with $(\bar{\epsilon}_i,\bar{\delta}_i)=(\ln 3, 0.2)$. For the cost, we choose $Q_{ij}=100$ for all i and j, and we set $R_{ii}=100$ and $R_{ij}=5$.

The effects of privacy on cost are shown in Figure 1. As expected, the cost with privacy is higher than without privacy. However, this increase becomes relatively modest over time, which indicates that privacy is well-suited to the long-horizon problems we consider. In Figure 2, we show the instantaneous

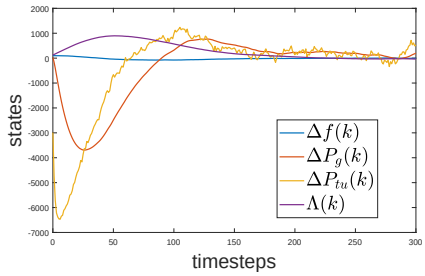

Fig. 3: Even with privacy, control values provided by the cloud are able to regulate an agent's state to remain near its desired trajectory.

error of the cloud's state estimates, and we compare that with the bounds in Theorem 1; we note that we plot the instantaneous error, but the bounds are for mean-square error. As expected, there are ephemeral bound violations by the instantaneous error, and it is shown that on average, the *a posteriori* error lies within the bounds in Theorem 1. This illustrates that privacy is compatible with the cloud estimating agents' states under privacy. Finally, Figure 3 illustrates the behaviour of the states of one of the areas.

IX. Conclusions

We have studied distributed linear-quadratic control with differential privacy and bounded the uncertainty and cost induced by privacy. Future work will develop differential privacy for other optimal control problems, including in model-free contexts at the intersection of control and learning [31].

References

[1] P. McDaniel and S. McLaughlin, "Security and privacy challenges in the smart grid," *IEEE Security Privacy*, vol. 7, no. 3, pp. 75–77, May 2009.

[2] U. D. of Energy, "Department of energy data access and privacy issues related to smart grid technologies," U.S. DoE, Tech. Rep., 2010.

[3] C. Dwork, "Differential privacy," in *33rd International Colloquium on Automata, Languages and Programming, part II (ICALP 2006)*, vol. 4052. Venice, Italy: Springer Verlag, July 2006, pp. 1–12.

[4] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proceedings of the Third Conference on Theory of Cryptography*. Springer-Verlag, 2006, pp. 265–284.

[5] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[6] J. L. Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, Feb 2014.

[7] J. L. Ny and M. Mohammady, "Differentially private mimo filtering for event streams and spatio-temporal monitoring," in *53rd IEEE Conference on Decision and Control*, Dec 2014, pp. 2148–2153.

[8] S. P. Kasiviswanathan and A. Smith, "On the semantics of differential privacy: A bayesian formulation," *arXiv preprint arXiv:0803.3946*, 2008.

[9] J. Le Ny, *Differential Privacy for Dynamic Data*. Springer, 2020.

[10] V. Rostampour, R. M. Ferrari, A. M. Teixeira, and T. Keviczky, "Privatized distributed anomaly detection for large-scale nonlinear uncertain systems," *IEEE Transactions on Automatic Control*, 2020.

[11] A. Sargolzaei, K. K. Yen, and M. N. Abdelghani, "Preventing time-delay switch attack on load frequency control in distributed power systems," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 1176–1185, 2016.

[12] L. Jiang, W. Yao, Q. H. Wu, J. Y. Wen, and S. J. Cheng, "Delay-dependent stability for load frequency control with constant and time-varying delays," *IEEE Transactions on Power Systems*, vol. 27, no. 2, pp. 932–941, 2012.

[13] D. Soudbakhsh, A. Chakrabortty, and A. M. Annaswamy, "A delay-aware cyber-physical architecture for wide-area control of power systems," *Control Engineering Practice*, vol. 60, pp. 171 – 182, 2017.

[14] D. Soudbakhsh, A. Chakrabortty, and A. M. Annaswamy, "Delay-aware co-designs for wide-area control of power grids," in *53rd IEEE Conference on Decision and Control*, 2014, pp. 2493–2498.

[15] A. K. Singh and B. C. Pal, "Decentralized control of oscillatory dynamics in power systems using an extended lqr," *IEEE Transactions on Power Systems*, vol. 31, no. 3, pp. 1715–1728, 2016.

[16] V. Katewa, A. Chakrabortty, and V. Gupta, "Differential privacy for network identification," *IEEE Transactions on Control of Network Systems*, vol. 7, no. 1, pp. 266–277, 2019.

[17] J. Chen, W. Zhan, and M. Tomizuka, "Autonomous driving motion planning with constrained iterative lqr," *IEEE Transactions on Intelligent Vehicles*, vol. 4, no. 2, pp. 244–254, 2019.

[18] H. Zhang and W. X. Zheng, "Denial-of-service power dispatch against linear quadratic control via a fading channel," *IEEE Transactions on Automatic Control*, vol. 63, no. 9, pp. 3032–3039, 2018.

[19] M. Hale and M. Egerstedt, "Cloud-based optimization: A quasi-decentralized approach to multi-agent coordination," in *53rd IEEE Conference on Decision and Control*, Dec 2014, pp. 6635–6640.

[20] M. Hale and M. Egerstedt, "Cloud-enabled differentially private multi-agent optimization with constraints," *IEEE Transactions on Control of Network Systems*, 2017, in press.

[21] M. Hale, A. Jones, and K. Leahy, "Privacy in feedback: The differentially private lqg," in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 3386–3391.

[22] K. Yazdani and M. Hale, "Error bounds and guidelines for privacy calibration in differentially private kalman filtering," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 4423–4428.

[23] J. J. Shynk, *Probability, random variables, and random processes: theory and signal processing applications*. John Wiley & Sons, 2012.

[24] D. P. Bertsekas, *Dynamic programming and optimal control*, 3rd ed. Athena Scientific Belmont, MA, 2005, vol. 1.

[25] B. Anderson and J. Moore, *Optimal Control: Linear Quadratic Methods*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1990.

[26] E. Sontag, *Mathematical control theory: deterministic finite dimensional systems*. Springer Science & Business Media, 2013, vol. 6.

[27] A. Sargolzaei, K. K. Yen, M. N. Abdelghani, S. Sargolzaei, and B. Carbunar, "Resilient design of networked control systems under time delay switch attacks, application in smart grid," *IEEE Access*, vol. 5, pp. 15 901–15 912, 2017.

[28] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.

[29] K. Yazdani and M. T. Hale, "Infinite horizon linear quadratic gaussian tracking control derivation," 2017, available at: https://arxiv.org/abs/1807.04700.

[30] J. Hsu, M. Gaboardi, A. Haeberlen, S. Khanna, A. Narayan, B. C. Pierce, and A. Roth, "Differential privacy: An economic method for choosing epsilon," in *Proceedings of the 2014 IEEE 27th Computer Security Foundations Symposium*, ser. CSF '14. Washington, DC, USA: IEEE Computer Society, 2014, pp. 398–410.

[31] F. L. Lewis, D. Vrabie, and K. G. Vamvoudakis, "Reinforcement learning and feedback control: Using natural decision methods to design optimal adaptive controllers," *IEEE Control Systems Magazine*, vol. 32, no. 6, pp. 76–105, 2012.

[32] D. S. Bernstein, *Matrix Mathematics: Theory, Facts, and Formulas: Second Edition*, 2nd ed. Princeton University Press, 2009.

[33] J. Garloff, "Bounds for the eigenvalues of the solution of the discrete riccati and lyapunov equations and the continuous lyapunov equation," *International Journal of Control*, vol. 43, no. 2, pp. 423–431, 1986.

[34] K. B. Petersen and M. S. Pedersen, "The matrix cookbook," nov 2012.

[35] A. Johnson, "Discrete and sampled-data stochastic control problems with complete and incomplete state information," *Applied Mathematics and Optimization*, vol. 24, no. 1, pp. 289–316, Jul 1991.

X. Appendix

The following lemmas will be used in deriving error bounds.

Lemma 2. [32, Fact 5.12.4] Let Υ and Θ be symmetric $n \times n$ matrices. If $\Upsilon \succ 0$, then $\lambda_n(\Theta)\mathrm{tr}(\Upsilon) \leq \mathrm{tr}(\Upsilon\Theta) \leq \lambda_1(\Theta)\mathrm{tr}(\Upsilon)$.

Lemma 3. [32, Theorem 8.4.11.] Let Υ and Θ be $n \times n$ Hermitian matrices. Then

$$\lambda_1(\Upsilon) + \lambda_n(\Theta) \leq \lambda_1(\Upsilon + \Theta) \leq \lambda_1(\Upsilon) + \lambda_1(\Theta)$$
$$\lambda_n(\Upsilon) + \lambda_n(\Theta) \leq \lambda_n(\Upsilon + \Theta) \leq \lambda_n(\Upsilon) + \lambda_1(\Theta).$$

Proof of Theorem 1: The steady-state MSE of the Kalman filter's predictions is $\text{tr}(\Sigma)$. Taking the trace of the Riccati equation defining Σ, we obtain $\text{tr}\Sigma - \text{tr}W = \text{tr}\Big[A^T A(\Sigma^{-1} + C^T V^{-1} C)^{-1}\Big]$, where we have used the cyclic permutation property of the trace. Next, we use Lemma 2 to write

$$\text{tr}\Sigma - \text{tr}W \geq \text{tr}(A^T A)\lambda_n\Big[(\Sigma^{-1} + C^T V^{-1} C)^{-1}\Big]$$
$$\geq \frac{\text{tr}(A^T A)}{\lambda_1(\Sigma^{-1}) + \lambda_1\big(C^T V^{-1} C\big)} = \frac{\text{tr}(A^T A)}{\frac{1}{\lambda_n(\Sigma)} + \lambda_1\big(C^T V^{-1} C\big)},$$

where we apply Lemma 3 on the second line to split up the eigenvalues and use the fact that $\lambda_1(\Sigma^{-1}) = {}^1\!/_{\lambda_n(\Sigma)}$ in the final step. It is shown in [33, Theorem 3.1] that $\Sigma \succeq W$, and therefore $\lambda_n(\Sigma) \geq \lambda_n(W)$. Using this fact and Equation (V) completes the first part of the proof. Similarly, applying Lemmas 2 and 3 to the same Riccati equation,

$$\text{tr}\Sigma - \text{tr}W \leq \text{tr}(A^T A)\lambda_1\Big[(\Sigma^{-1} + C^T V^{-1} C)^{-1}\Big]$$
$$\leq \frac{\text{tr}(A^T A)}{\lambda_n(\Sigma^{-1}) + \lambda_n(C^T V^{-1} C)} \leq \frac{\sigma_l^2 \text{tr}(A^T A)}{C_l^2},$$

where the second step uses $\lambda_1(\Upsilon^{-1}) = {}^1\!/_{\lambda_n(\Upsilon)}$ and the third step uses Lemma 3 to split the eigenvalues.

The steady-state MSE of the Kalman filter's state estimates is $\text{tr}(\overline{\Sigma})$. Using Lemma 2,

$$\text{tr}\overline{\Sigma} \geq \frac{n}{\lambda_1(C^T V^{-1} C + \Sigma^{-1})} \geq \frac{n}{\lambda_1(C^T V^{-1} C) + \lambda_1(\Sigma^{-1})}$$
$$\geq \frac{n}{\lambda_1(C^T V^{-1} C) + \lambda_n^{-1}(W)} = \frac{n\sigma_u^2}{C_u^2 + \sigma_u^2 \lambda_n^{-1}(W)},$$

where in the second inequality we use Lemma 3 to split the eigenvalues. In the last line, we use $\lambda_n(\Sigma) \geq \lambda_n(W)$ [33, Theorem 3.1] and Equation (V).

Using Lemma 2, an upper bound can be derived with

$$\text{tr}\overline{\Sigma} \leq n\lambda_1\big((C^T V^{-1} C + \Sigma^{-1})^{-1}\big) \leq \frac{n}{\lambda_n(C^T V^{-1} C)},$$

where we have used Lemma 3 to split the eigenvalues. Using Equation (V) completes the proof. ■

Proof of Theorem 2: Choose $\epsilon_i \geq \frac{1}{8}\left(\frac{1+\sqrt{36\eta_4+1}}{\eta_4}\right)^2$ and solve for η_4 to get $\frac{9+\sqrt{2\epsilon_i}}{2\epsilon_i} \leq \eta_4$. Choosing $\delta_i \in [10^{-5}, 10^{-1}]$ gives $K_{\delta_i} \in [1, 4.5]$. Then $\frac{2K_{\delta_i}+\sqrt{2\epsilon_i}}{2\epsilon_i} \leq \eta_4$. Because $\sqrt{\upsilon+\theta} \leq \sqrt{\upsilon} + \sqrt{\theta}$, we can lower-bound the left-hand-side to write $\kappa(\delta_i, \epsilon_i) \leq \eta_4$. Squaring, substituting in η_4, and rearranging we find $s_1(C_i)^2 \kappa(\delta_i, \epsilon_i)^2 b_i^2 \leq \frac{B_u C_l^2}{n}$, which implies $\sigma_l^2 \leq \frac{B_u C_l^2}{n}$. Comparing to Theorem 1, we see that $\text{tr}\overline{\Sigma} \leq B_u$.

Next, choose $\epsilon_i \leq \frac{1}{\eta_3}$. Given $K_{\delta_i} \in [1, 4.5]$, we may write $\eta_3 \leq \frac{K_{\delta_i}}{\epsilon_i}$. We substitute for η_3 and square both sides to write

$$\frac{B_l C_u^2}{s_1(C_i)^2 b_i^2 \left(n - B_l \lambda_n^{-1}(W)\right)} \leq \left(\frac{K_{\delta_i}}{\epsilon_i}\right)^2.$$

Using $\frac{K_{\delta_i}}{\epsilon_i} \leq \kappa(\delta_i, \epsilon_i)$ and rearranging, we have

$$\frac{B_l C_u^2}{n - B_l \lambda_n^{-1}(W)} \leq s_1(C_i)^2 \kappa(\delta_i, \epsilon_i)^2 b_i^2.$$

This implies $\frac{B_l C_u^2}{n - B_l \lambda_n^{-1}(W)} \leq \sigma_u^2$. Isolating B_l and applying Theorem 1 implies $\text{tr}\overline{\Sigma} \geq B_l$. ■

Proof of Theorem 3: The cost of privatizing $\bar{x}$ specifically is

$$\text{tr}\big([Q + H^T R H]\overline{W}\big), \tag{5}$$

which is obtained from the authors' technical report in [29] and application of [34, Equation (318)]. Next, using [35] and Equation (X), the total cost incurred by Algorithm 1 is

$$\tilde{J}(x,u) = J(x,u) + \text{tr}(Q\overline{W}) + \text{tr}(H^T R H \overline{W}) =$$
$$\lim_{K_f \to \infty} \frac{1}{K_f} \sum_{k=1}^{K_f} \text{tr}\left(K\Sigma + (Q-K)\,\overline{\Sigma}\right)$$
$$+ \lim_{K_f \to \infty} \frac{1}{K_f} \sum_{k=0}^{K_f - 1} \bar{x}^T Q \bar{x} - g^T B \left(R + B^T K B\right)^{-1} B^T g$$
$$+ \text{tr}(Q\overline{W}) + \text{tr}(H^T R H \overline{W}),$$

where the second step follows from [35, Equation 4.12]. Then (3) follows by subtracting the cost without privacy noise. ■

Proof of Theorem 4: Using chain rule we have $\frac{d\Delta J}{d\epsilon} = \frac{d\Delta J}{d\sigma}\frac{d\sigma}{d\epsilon}$. For the first term, we have

$$\frac{d\Delta J(x,u)}{d\sigma} = \text{tr}\left[\frac{d\left(K\Sigma + (Q-K)\,\overline{\Sigma}\right)}{d\sigma}\right]$$
$$+ \frac{d\left[-\text{tr}\,(KW) + \text{tr}(Q\overline{W}) + \text{tr}(H^T R H \overline{W})\right]}{d\sigma} \tag{6}$$

where we have used [34, Equation 36] to move the derivative inside the trace. The matrices K and $Q - K$ are symmetric, and $\text{tr}\frac{d\Sigma}{d\sigma}, \text{tr}\frac{d\overline{\Sigma}}{d\sigma} > 0$ because filter error monotonically increases with privacy noise. Therefore, by Lemma 2, the first term in (X) can be bounded by

$$\lambda_n(K)\text{tr}\frac{d\Sigma}{d\sigma} + \lambda_n(Q-K)\text{tr}\frac{d\overline{\Sigma}}{d\sigma} \leq \text{tr}\frac{d\left(K\Sigma + (Q-K)\,\overline{\Sigma}\right)}{d\sigma}$$
$$\leq \lambda_1(K)\text{tr}\frac{d\Sigma}{d\sigma} + \lambda_1(Q-K)\text{tr}\frac{d\overline{\Sigma}}{d\sigma}. \tag{7}$$

By differentiating the discrete algebraic Riccati equation that defines Σ, we have

$$\frac{d\Sigma}{d\sigma} = A\frac{d\Sigma}{d\sigma}A^T - A\frac{d\Sigma}{d\sigma}C^T\left(C\Sigma C^T + V\right)^{-1} C\Sigma A^T$$
$$- A\Sigma C^T\left(C\Sigma C^T + V\right)^{-1} C\frac{d\Sigma}{d\sigma}A^T$$
$$+ A\Sigma C^T\big(C\Sigma C^T + V\big)^{-1}\left(C\frac{d\Sigma}{d\sigma}C^T + 2\sigma I\right)\big(C\Sigma C^T + V\big)^{-1} C\Sigma A^T.$$

Taking the trace of both sides and simplifying, we find $\text{tr}\left(U\frac{d\Sigma}{d\sigma}\right) = -2\sigma\text{tr}(F^T F)$. The matrix U is symmetric, and, because A is stable, it is positive definite. Therefore by applying Lemma 2 and simplifying we find

$$\max\left\{\frac{-2\sigma\text{tr}(F^T F)}{\lambda_1(U)}, 0\right\} \leq \text{tr}\frac{d\Sigma}{d\sigma} \leq \frac{-2\sigma\text{tr}(F^T F)}{\lambda_n(U)}. \tag{8}$$

Next we differentiate the equation defining $\overline{\Sigma}$ to find

$$\frac{d\overline{\Sigma}}{d\sigma} = \frac{d\Sigma}{d\sigma} - \frac{d\Sigma}{d\sigma}C^T\left(C\Sigma C^T + V\right)^{-1}C\Sigma - \Sigma C^T\left(C\Sigma C^T + V\right)^{-1}C\frac{d\Sigma}{d\sigma} + \Sigma C^T(C\Sigma C^T+V)^{-1}\left(C\frac{d\Sigma}{d\sigma}C^T+2\sigma I\right)(C\Sigma C^T+V)^{-1}C\Sigma.$$

Taking the trace gives $\text{tr}\frac{d\overline{\Sigma}}{d\sigma} = \text{tr}\left[\frac{d\Sigma}{d\sigma}\bar{U}\right] + \text{tr}\left(2\sigma\bar{F}^T\bar{F}\right)$. By substituting the bounds in Equation (X) we get

$$\max\left\{\frac{-2\sigma\text{tr}(F^TF)}{\lambda_1(U)},0\right\}\lambda_n\left(\bar{U}\right) + \text{tr}\left(2\sigma\bar{F}^T\bar{F}\right) \leq \text{tr}\frac{d\overline{\Sigma}}{d\sigma} \leq \frac{-2\sigma\text{tr}(F^TF)}{\lambda_n(U)}\lambda_1\left(\bar{U}\right) + \text{tr}\left(2\sigma\bar{F}^T\bar{F}\right). \quad (9)$$

Substituting the results from Equations (X) and (X) into (X) and assemble the results back in Equation (X), we get

$$\lambda_n(K)\max\left\{\frac{-2\sigma\text{tr}(F^TF)}{\lambda_1(U)},0\right\} + \lambda_n(Q-K)\left[\frac{-2\sigma\text{tr}(F^TF)}{\lambda_n(U)}\lambda_1\left(\bar{U}\right)+\text{tr}\left(2\sigma\bar{F}^T\bar{F}\right)\right]+2\sigma\text{tr}Q+2\sigma\text{tr}\left(H^TRH\right) \leq \frac{d\Delta J}{d\sigma} \leq \lambda_1(K)\frac{-2\sigma\text{tr}(F^TF)}{\lambda_n(U)} + \lambda_1(Q-K)\left[\max\left\{\frac{-2\sigma\text{tr}(F^TF)}{\lambda_1(U)},0\right\}\lambda_n\left(\bar{U}\right)+\text{tr}\left(2\sigma\bar{F}^T\bar{F}\right)\right] + 2\sigma\text{tr}Q + 2\sigma\text{tr}\left(H^TRH\right). \quad (10)$$

Next, we observe that $\frac{d\sigma}{d\epsilon} < 0$. We multiply Equation (X) by $\frac{d\sigma}{d\epsilon} < 0$ and this completes the proof. ■

Proof of Theorem 5: Choosing $\epsilon_i \geq \frac{1}{8}\left(\frac{1+\sqrt{36\eta_5+1}}{\eta_5}\right)^2$ and solving for η_5, we find $\frac{9+\sqrt{2\epsilon_i}}{2\epsilon_i} \leq \eta_5$. Taking $\delta \in [10^{-5}, 10^{-1}]$ implies $K_{\delta_i} \in [1, 4.5]$, and as a result we can write $\frac{2K_{\delta_i}+\sqrt{2\epsilon_i}}{2\epsilon_i} \leq \eta_5$. Using $\sqrt{v+\theta} \leq \sqrt{v} + \sqrt{\theta}$ we take K_{δ_i} inside the square root, leading to $\kappa(\delta_i, \epsilon_i) \leq \eta_5$. Expanding, this is equivalent to

$$\sigma_i^2 \leq \frac{B_u - \lambda_1(K)\text{tr}W - \bar{x}^TQ\bar{x} + g^TB\left(R+B^TKB\right)^{-1}B^Tg}{\lambda_1\left(K\right)\frac{\text{tr}(A^TA)}{C_l^2} + \text{tr}(H^TRH) + \text{tr}(Q)}.$$

By rearranging terms and using $\bar{\sigma}_i = \sigma_i = \sigma_l$, we find

$$\lambda_1\left(K\right)\left[\text{tr}W + \frac{\sigma_l^2\text{tr}(A^TA)}{C_l^2}\right] + \bar{x}^TQ\bar{x} - g^TB\left(R+B^TKB\right)^{-1}B^Tg + \bar{\sigma}_i^2\left[\text{tr}(H^TRH) + \text{tr}(Q)\right] \leq B_u.$$

Using $\bar{\sigma}_i^2 \leq \lambda_1(\overline{W})$ and Lemma 2, we have $\text{tr}(Q\overline{W}) + \text{tr}(H^TRH\overline{W}) \leq \text{tr}(Q)\lambda_1\left(\overline{W}\right) + \text{tr}(H^TRH)\lambda_1\left(\overline{W}\right)$ and we can write

$$\lambda_1\left(K\right)\left[\text{tr}W + \frac{\sigma_l^2\text{tr}(A^TA)}{C_l^2}\right] + \bar{x}^TQ\bar{x} - g^TB\left(R+B^TKB\right)^{-1}B^Tg + \text{tr}(Q\overline{W}) + \text{tr}(H^TRH\overline{W}) \leq B_u. \quad (11)$$

From [33, Theorem 3.1] we know that $Q - K \preceq 0$ and thus $\lambda_1(Q-K) \leq 0$. Using this in Equation (X), we find

$$\lambda_1\left(K\right)\left[\text{tr}W + \frac{\sigma_l^2\text{tr}(A^TA)}{C_l^2}\right] + \lambda_1\left(Q-K\right)\left[\frac{n\sigma_l^2}{C_l^2+\sigma_l^2\lambda_n^{-1}(W)}\right] + \bar{x}^TQ\bar{x} - g^TB\left(R+B^TKB\right)^{-1}B^Tg+\text{tr}(Q\overline{W})+\text{tr}(H^TRH\overline{W}) \leq B_u.$$

Using Theorem 1, we can write

$$\lambda_1\left(K\right)\text{tr}\left(\Sigma\right) + \lambda_1\left(Q-K\right)\text{tr}\overline{\Sigma} + \bar{x}^TQ\bar{x} - g^TB\left(R+B^TKB\right)^{-1}B^Tg+\text{tr}(Q\overline{W})+\text{tr}(H^TRH\overline{W}) \leq B_u,$$

and therefore

$$\text{tr}\left(K\Sigma\right) + \text{tr}\left[\left(Q-K\right)\overline{\Sigma}\right] + \bar{x}^TQ\bar{x} - g^TB\left(R+B^TKB\right)^{-1}B^Tg+\text{tr}(Q\overline{W})+\text{tr}(H^TRH\overline{W}) \leq B_u,$$

and we find $\tilde{J}\left(x,u\right) \leq \alpha$, which completes the proof. ■